Freeform Search

Database: US Pre-Grant Publication Full-Text Database; US Patents Full-Text Database; US OCR Full-Text Database; EPO Abstracts Database; JPO Abstracts Database; Derwent World Patents Index; IBM Technical Disclosure Bulletins

Term:

Display: 10 **Documents in Display Format:** - **Starting with Number** 1

Generate: ○ Hit List ◉ Hit Count ○ Side by Side ○ Image

Search | Clear | Interrupt

Search History

DATE: Sunday, December 11, 2005 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=PGPB,USPT; PLUR=YES; OP=OR			
L13	l3 and (((usage or used or repeated or times) near4 (writ$3 or read$3 or transfer$4)) near6 (large or big or threshold))	124	L13
DB=PGPB; PLUR=YES; OP=OR			
L12	L11 and (cache or logical or redundan$4)	1	L12
L11	(20040250021).pn.	1	L11
DB=PGPB,USPT; PLUR=YES; OP=OR			
L10	l9 and ((error or fail$4 or fault) near8 (transfer$4 or writ43 or read$3))	21	L10
L9	L7 and (redundan$3 or mirror$3)	32	L9
L8	L7 and redundan$3	32	L8
L7	L6 and (network or LAN or WAN)	35	L7
L6	L5 and cache	66	L6
L5	L4 and l3	113	L5
L4	(714/5,1,2,25,42,47,100).ccls.	4170	L4
L3	(711/113,114).ccls.	2719	L3
L2	l1 and (error or fail$4)	2	L2
L1	(((5761411 or 6061761 or 6366981 or 6611896 or 6715054 or 20010054133).pn.) and cache)	4	L1

END OF SEARCH HISTORY

Freeform Search

Database: US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Term:

Display: 10 **Documents in Display Format:** TI **Starting with Number** 1

Generate: ○ **Hit List** ◉ **Hit Count** ○ **Side by Side** ○ **Image**

Search | Clear | Interrupt

Search History

DATE: Sunday, December 11, 2005 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=PGPB,USPT; PLUR=YES; OP=OR			
L18	(RAID or array or disk or disc) and (((queue or buffer) near3 number) near8 redundan$4 near8 (compar$4 or examin$6 or relat$3))	4	L18
DB=USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
L17	(RAID or array or disk or disc) and (((queue or buffer) near3 number) near8 redundan$4 near8 (compar$4 or examin$6 or relat$3))	0	L17
L16	(RAID or array or disk or disc) and (((average or ratio) near3 time) near8 redundan$4 near8 (compar$4 or examin$6 or relat$3))	0	L16
DB=PGPB,USPT; PLUR=YES; OP=OR			
L15	(RAID or array or disk or disc) and (((average or ratio) near3 time) near8 redundan$4 near8 (compar$4 or examin$6 or relat$3))	8	L15
L14	l13 and (((average or ratio) near3 time) near8 redundan$4)	2	L14
L13	l3 and (((usage or used or repeated or times) near4 (writ$3 or read$3 or transfer$4)) near6 (large or big or threshold))	124	L13
DB=PGPB; PLUR=YES; OP=OR			
L12	L11 and (cache or logical or redundan$4)	1	L12
L11	(20040250021).pn.	1	L11
DB=PGPB,USPT; PLUR=YES; OP=OR			
L10	l9 and ((error or fail$4 or fault) near8 (transfer$4 or writ43 or read$3))	21	L10
L9	L7 and (redundan$3 or mirror$3)	32	L9
L8	L7 and redundan$3	32	L8
L7	L6 and (network or LAN or WAN)	35	L7

L6	L5 and cache	66	L6
L5	L4 and l3	113	L5
L4	(714/5,1,2,25,42,47,100).ccls.	4170	L4
L3	(711/113,114).ccls.	2719	L3
L2	l1 and (error or fail$4)	2	L2
L1	(((5761411 or 6061761 or 6366981 or 6611896 or 6715054 or 20010054133).pn.) and cache)	4	L1

END OF SEARCH HISTORY

Results for "(((disk or array or raid) and (((average or ratio) <near/3> time) <near/8> redundan$ <..."

e-mail

Your search matched **0** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

(((disk or array or raid) and (((average or ratio) <near/3> time) <near/8> redundan$ <

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan search.

Results for "(((raid or array or disk or disc) and (((queue or buffer) <near/3> number) <near/8> redu..."

e-mail

Your search matched **0** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

(((raid or array or disk or disc) and (((queue or buffer) <near/3> number) <near/8> red

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan search.

Nothing Found

Your search for **(RAID or array or disk or disc) and (((queue or buffer) <near/3> number) <near/8> redundan* <near/8> (compar* or examin* or relat*)) <near/8> redundan* <near/8> (compar* or examin* or relat*))** did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

- Enter your search terms in lower case with a space between the terms.

  ```
  sales offices
  ```

  You can also enter a full question or concept in plain language.

  ```
  Where are the sales offices?
  ```

- Capitalize proper nouns to search for specific people, places, or products.

  ```
  John Colter, Netscape Navigator
  ```

- Enclose a phrase in double quotes to search for that exact phrase.

  ```
  "museum of natural history" "museum of modern art"
  ```

- Narrow your searches by using a + if a search term must appear on a page.

  ```
  museum +art
  ```

- Exclude pages by using a - if a search term must not appear on a page.

  ```
  museum -Paris
  ```

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

```
museum +"natural history" dinosaur -Chicago
```

1 Parity declustering for continuous operation in redundant disk arrays

Mark Holland, Garth A. Gibson

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems ASPLOS-V**, Volume 27 Issue 9

Publisher: ACM Press

Full text available: pdf(1.57 MB) Additional Information: full citation, references, citings, index terms

2 Parity logging disk arrays

Daniel Stodolsky, Mark Holland, William V. Courtright, Garth A. Gibson

August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3

Publisher: ACM Press

Full text available: pdf(1.98 MB) Additional Information: full citation, abstract, references, citings, index terms

Parity-encoded redundant disk arrays provide highly reliable, cost-effective secondary storage with high performance for reads and large writes. Their performance on small writes, however, is much worse than mirrored disks—the traditional, highly reliable, but expensive organization for secondary storage. Unfortunately, small writes are a substantial portion of the I/O workload of many important, demanding applications such as on-line transaction processing. This paper presents

Keywords: RAID, disk arrays

3 Destage algorithms for disk arrays with non-volatile caches

Anujan Varma, Quinn Jacobson

May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture ISCA '95**, Volume 23 Issue 2

Publisher: ACM Press

Full text available: pdf(1.63 MB) Additional Information: full citation, abstract, references, citings, index terms

In a disk array with a nonvolatile write cache, destages from the cache to the disk are performed in the background asynchronously while read requests from the host system are serviced in the foreground. In this paper, we study a number of algorithms for scheduling destages in a RAID-5 system. We introduce a new scheduling algorithm, called *linear threshold scheduling,* that adaptively varies the rate of destages to disks based on the instantaneous occupancy of the write cache. The perform ...

4 Parity logging overcoming the small write problem in redundant disk arrays

Daniel Stodolsky, Garth Gibson, Mark Holland

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: pdf(1.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Parity encoded redundant disk arrays provide highly reliable, cost effective secondary storage with high performance for read accesses and large write accesses. Their performance on small writes, however, is much worse than mirrored disks—the traditional, highly reliable, but expensive organization for secondary storage. Unfortunately, small writes are a substantial portion of the I/O workload of many important, demanding applications such as on-line transaction processing. This paper ...

5 Deconstructing storage arrays

Timothy E. Denehy, John Bent, Florentina I. Popovici, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

October 2004 **ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 39 , 32 , 38 Issue 11 , 5 , 5

Publisher: ACM Press

Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, index terms

We introduce Shear, a user-level software tool that characterizes RAID storage arrays. Shear employs a set of controlled algorithms combined with statistical techniques to automatically determine the important properties of a RAID system, including the number of disks, chunk size, level of redundancy, and layout scheme. We illustrate the correctness of Shear by running it upon numerous simulated configurations, and then verify its real-world applicability by running Shear on both software-based ...

Keywords: RAID, storage

6 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson

June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

7 Declustered disk array architectures with optimal and near-optimal parallelism

Guillermo A. Alvarez, Walter A. Burkhard, Larry J. Stockmeyer, Flaviu Cristian

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.35 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

This paper investigates the placement of data and parity on redundant disk arrays. Declustered organizations have been traditionally used to achieve fast reconstruction of a failed disk's contents. In previous work, Holland and Gibson identified six desirable properties for ideal layouts; however, no declustered layout satisfying all properties has been published in the literature. We present a complete, constructive characterization of the collection of ideal declustered layouts possessing all ...

8 Hot mirroring: a method of hiding parity update penalty and degradation during rebuilds for RAID5

Kazuhiko Mogi, Masaru Kitsuregawa

June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international**

conference on Management of data SIGMOD '96, Volume 25 Issue 2
Publisher: ACM Press
Full text available: pdf(1.37 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper proposes a storage management scheme for disk arrays, named hot mirroring. In this scheme, storage space is partitioned into two regions. One is the mirrored region, which is characterized by high performance and low storage efficiency. The other is the RAID5 region, which is characterized by low performance and high storage efficiency. Hot data blocks are stored in the former area, while cold blocks are stored in the latter. In addition, mirrored pairs and RAID5 stripes are orthogona ...

9 An analytic performance model of disk arrays
Edward K. Lee, Randy H. Katz
June 1993 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '93**, Volume 21 Issue 1
Publisher: ACM Press
Full text available: pdf(1.13 MB) Additional Information: full citation, abstract, references, citings, index terms

As disk arrays become widely used, tools for understanding and analyzing their performance become increasingly important. In particular, performance models can be invaluable in both configuring and designing disk arrays. Accurate analytic performance models are preferable to other types of models because they can be quickly evaluated, are applicable under a wide range of system and workload parameters, and can be manipulated by a range of mathematical techniques. Unfortunately, analytic performa ...

10 Disk array performance in a random IO environment
T. M. Oslon
September 1989 **ACM SIGARCH Computer Architecture News**, Volume 17 Issue 5
Publisher: ACM Press
Full text available: pdf(441.34 KB) Additional Information: full citation, abstract, citings, index terms

Large arrays of disks have been proposed as a way to meet the need for increasing IO bandwidth. This paper examines disk array performance in a random IO environment. It also presents the results of performance testing using the Prime IOBENCH™ benchmark on a combination of disk striping, RAID 1, and RAID 5 disk arrays. It concludes, that for a given number of disk drives, that both RAID 1 and RAID 5 have acceptable performance in a read environment, while RAID 5 degrades significantly ...

11 The TickerTAIP parallel RAID architecture
Pei Cao, Swee Boon Lin, Shivakumar Venkataraman, John Wilkes
August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3
Publisher: ACM Press
Full text available: pdf(2.04 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum number of disks to which the array can scale. We describe TickerTAIP, a parallel architecture for disk arrays that distributes the controller functions across several loosely coupled processors. The result is better scalability, fault tolerance, and flexibility.This article present ...

Keywords: RAID disk array, decentralized parity calculation, disk scheduling, distributed controller, fault tolerance, parallel controller, performance simulation

12 Comparison of sparing alternatives for disk arrays
Jai Menon, Dick Mattson
April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture ISCA '92**, Volume 20 Issue 2
Publisher: ACM Press
Full text available: pdf(1.27 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper explores how choice of sparing methods impacts the performance of RAID level 5

(or parity striped) disk arrays. The three sparing methods examined are dedicated sparing, distributed sparing, and parity sparing. For database type workloads with random single block reads and writes, array performance is compared in four different modes - normal mode (no disks have failed), degraded mode (a disk has failed and its data has not been reconstructed), rebuild mode (a disk has failed and ...

13 Failure correction techniques for large disk arrays

G. A. Gibson, L. Hellerstein, R. M. Karp, D. A. Patterson

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the third international conference on Architectural support for programming languages and operating systems ASPLOS-III**, Volume 17 Issue 2

Publisher: ACM Press

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms

The ever increasing need for I/O bandwidth will be met with ever larger arrays of disks. These arrays require redundancy to protect against data loss. This paper examines alternative choices for encodings, or codes, that reliably store information in disk arrays. Codes are selected to maximize mean time to data loss or minimize disks containing redundant data, but are all constrained to minimize performance penalties associated with updating information or recovering from catastroph ...

14 FAB: building distributed enterprise disk arrays from commodity components

Yasushi Saito, Svend Frølund, Alistair Veitch, Arif Merchant, Susan Spence

October 2004 **ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 32 , 38 , 39 Issue 5 , 5 , 11

Publisher: ACM Press

Full text available: pdf(671.67 KB) Additional Information: full citation, abstract, references, index terms

This paper describes the design, implementation, and evaluation of a Federated Array of Bricks (FAB), a distributed disk array that provides the reliability of traditional enterprise arrays with lower cost and better scalability. FAB is built from a collection of *bricks*, small storage appliances containing commodity disks, CPU, NVRAM, and network interface cards. FAB deploys a new majority-voting-based algorithm to replicate or erasure-code logical blocks across bricks and a reconfigurati ...

Keywords: consensus, disk array, erasure coding, replication, storage, voting

15 X-RAY: A Non-Invasive Exclusive Caching Mechanism for RAIDs

Lakshmi N. Bairavasundaram, Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(250.59 KB) Additional Information: full citation, abstract, citings

RAID storage arrays often possess gigabytes of RAM forcaching disk blocks. Currently, most RAID systems use LRUor LRU-like policies to manage these caches. Since these arraycaches do not recognize the presence of file system buffer caches,they redundantly retain many of the same blocks as those cachedby the file system, thereby wasting precious cache space. In thispaper, we introduce X-RAY, an exclusive RAID array cachingmechanism. X-RAY achieves a high degree of (but not perfect) exclusivitythr ...

16 The TickerTAIP parallel RAID architecture

Pei Cao, Swee Boon Lim, Shivakumar Venkataraman, John Wilkes

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: pdf(1.19 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the

maximum size that the array can grow to. We describe here TickerTAIP, a parallel architecture for disk arrays that distributed the controller functions across several loosely-coupled processors. The result is better scalability, fault tolerance, and flexibility. This paper presents the Tic ...

17 Performance of a mirrored disk in a real-time transaction system

Shenze Chen, Don Towsley

April 1991 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1991 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '91**, Volume 19 Issue 1

Publisher: ACM Press

Full text available: pdf(975.06 KB) Additional Information: full citation, abstract, references, index terms

Disk mirroring has found widespread use in computer systems as a method for providing fault tolerance. In addition to increasing reliability, a mirrored disk can also reduce I/O response time by supporting the execution of parallel I/O requests. The improvement in I/O efficiency is extremely important in a real-time system, where each computational entity carries a deadline. In this paper, we present two classes of real-time disk scheduling policies, RT-DMQ and RT-CMQ, for a mirrored disk I/O su ...

18 External memory algorithms and data structures: dealing with **massive data**

Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Publisher: ACM Press

Full text available: pdf(828.46 KB) Additional Information: full citation, abstract, references, citings, index terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

19 Disk Drive Roadmap from the Thermal Perspective: A Case for Dynamic Thermal Management

Sudhanva Gurumurthi, Anand Sivasubramaniam, Vivek K. Natarajan

June 2005 **Proceedings of the 32nd Annual International Symposium on Computer Architecture ISCA '05**

Publisher: IEEE Computer Society

Full text available: pdf(243.57 KB) Additional Information: full citation, abstract

The importance of pushing the performance envelope of disk drives continues to grow, not just in the server market but also in numerous consumer electronics products. One of the most fundamental factors impacting disk drive design is the heat dissipation and its effect on drive reliability, since high temperatures can cause off-track errors, or even head crashes. Until now, drive manufacturers have continued to meet the 40% annual growth target of the internal data rates (IDR) by increasing RPMs ...

20 Disk cache—miss ratio analysis and design considerations

Alan J. Smith

August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Publisher: ACM Press

Full text available: pdf(3.13 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The current trend of computer system technology is toward CPUs with rapidly increasing processing power and toward disk drives of rapidly increasing density, but with disk performance increasing very slowly if at all. The implication of these trends is that at some point the processing power of computer systems will be limited by the throughput of the

input/output (I/O) system. A solution to this problem, which is described and evaluated in this paper, is disk cache